LUND UNIVERSITY

Towards Mission-Critical Control at the Edge and Over 5G

Skarin, Per; Tärneberg, William; Årzén, Karl-Erik; Kihl, Maria

Published in:
2018 IEEE International Conference on Edge Computing (EDGE)

DOI:
[10.1109/EDGE.2018.00014](https://doi.org/10.1109/EDGE.2018.00014)

2018

Document Version:
Publisher's PDF, also known as Version of record

[Link to publication](#)

Citation for published version (APA):
Skarin, P., Tärneberg, W., Årzén, K-E., & Kihl, M. (2018). Towards Mission-Critical Control at the Edge and Over 5G. In *2018 IEEE International Conference on Edge Computing (EDGE)* IEEE Computer Society.
<https://doi.org/10.1109/EDGE.2018.00014>

Total number of authors:
4

General rights

Unless other specific re-use rights are stated the following general rights apply:
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Towards Mission-Critical Control at the Edge and Over 5G

Per Skarin^{*†‡}, William Tärneberg^{*§}, Karl-Erik Årzen[†], and Maria Kihl[§]

^{*}These authors contributed equally to this work

[†] Department of Automatic Control, Lund University, Sweden

[‡]Ericsson Research, Lund, Sweden

[§]Department of Electrical and Information Technology, Lund University, Sweden

Abstract—

With the emergence of industrial IoT and cloud computing, and the advent of 5G and edge clouds, there are ambitious expectations on elasticity, economies of scale, and fast time to market for demanding use cases in the next generation of ICT networks. Responsiveness and reliability of wireless communication links and services in the cloud are set to improve significantly as the concept of edge clouds is becoming more prevalent. To enable industrial uptake we must provide cloud capacity in the networks but also a sufficient level of simplicity and self-sustainability in the software platforms. In this paper, we present a research test-bed built to study mission-critical control over the distributed edge cloud. We evaluate system properties using a conventional control application in the form of a Model Predictive Controller. Our cloud platform provides the means to continuously operate our mission-critical application while seamlessly relocating computations across geographically dispersed compute nodes. Through our use of 5G wireless radio, we allow for mobility and reliably provide compute resources with low latency, at the edge. The primary contribution of this paper is a state-of-the art, fully operational test-bed showing the potential for merged IoT, 5G, and cloud. We also provide an evaluation of the system while operating a mission-critical application and provide an outlook on a novel research direction.

Keywords—Edge cloud computing, Cloud computing, Time-sensitive, Mission-critical, Test-bed, 5G, IoT, Control theory, URLLC

I. INTRODUCTION

In this paper we target the feasibility of running time-sensitive and mission-critical applications in the edge cloud. We proceed by designing and building an edge cloud research test-bed that encompasses a distributed set of compute nodes, a distributed Platform-as-a-Service (PaaS) framework, a Fifth Generation Wireless Specifications (5G) cell, and a time-sensitive mission-critical process under control, see Figure 1.

A mission-critical system is one in which a failure or interruption comes with an unacceptable business or human cost. Here, a failure may be an arbitrarily small deviation from the desired operation. Naturally, this includes all systems that create a risk of injury but also systems in which failure incurs a very notable inconvenience, such as a means of transportation rendered practically useless because it is making passengers nauseous. Such applications are time-sensitive in the manner that they are unable to cope with delay and jitter in the delay, to the point where they violate their requirements.

The Wide Area Networks (WANs) separating a time-sensitive mission-critical system from a traditional distant Data Center (DC) may incur latencies beyond what is operationally

acceptable. The edge cloud was proposed to mitigate the latent latency, throughput, and availability barriers that separate the end users from distant DCs. When accessing the edge cloud over Ultra-Reliable and Low-Latency Communication (URLLC) 5G [1], the latencies are sufficiently low that time-sensitive mission-critical applications can be deployed in the edge cloud. An additional benefit of the edge cloud is that resident applications can be made spatially redundant and fall-back solutions can be implemented at various geographical points in the infrastructure for additional resilience.

Historically, control systems have been deployed as monolithic Software (SW) implementations on carefully tuned Hardware (HW), adjacent to the plants they control. Deploying monolithic SW on static HW makes such systems undesirably non-modular, less extensible, and limits their ability to self-adapt. Conversely, cloud-native applications are built for the cloud, offering the prospect of greater flexibility, reuse, availability, and reliability with lower latencies. When applications are implemented in a disaggregated manner, their execution can be distributed across the system's many nodes, migrated, and scaled to meet their individual objectives as well as that of the system as a whole. To adapt to, and prosper in, the edge cloud, applications will arguably have to adhere to a cloud-native paradigm.

The premise of this paper is that deploying mission-critical applications over the cloud, with wireless devices, must arguably occur in conjunction with the availability of edge cloud resources, the flexibility of cloud-native applications, and the reliability and low latency of 5G. We argue that such applications can operate in and make use of a distributed edge cloud but that there need to be relevant tools for them to be native to this context. There are many challenges and performance uncertainties in this premise. Therefore, we study the feasibility of deploying time-sensitive mission-critical applications and their performance when deployed on an actual edge cloud infrastructure. The contributions of this paper are:

- A state-of-the art edge cloud research test-bed aimed at the study of software autonomy and mission-critical applications.
- An empirical baseline evaluation of the plausibility of deploying latency-sensitive applications in the edge cloud.
- An empirical evaluation of the system's ability to dynamically reconfigure during run-time and the impact this has on the application.
- An empirical evaluation of the benefits of deploying

Fig. 1: System overview

latency-sensitive mission-critical applications at the edge, at the plant, and on a distant DC.

Section II covers the related work in the field and highlights the research gap. This is followed by a detailed account of the implemented test-bed in Section III. Section IV presents an example automatic control application which is used to evaluate the test-bed. Finally, Section V highlights the contributions of the paper and points to new research directions.

II. RELATED WORK

In this section, we survey related work and highlight the apparent research gaps left by the literature. The works below cover both related attempts at research test-beds and experiments that pursue the viability of the edge cloud.

Test-beds spanning User Equipments (UEs), wired and wireless networks, distributed cloud infrastructure and platforms are crucial instruments to realise and study the complexity of the edge cloud. The literature contains a number of such attempts. The authors of [2] present the SAVI test-bed which is an edge cloud test-bed realising Network Function Virtualisation (NFV) with a Field-Programmable Gate Array (FPGA)-cloud. SAVI is used in the investigation of virtualising the wireless access network. Although comprehensive, the test-bed does not provide a general edge cloud implementation for cloud native applications nor does the implementation span multiple tiers of cloud resources, including the device. In [3], a full test-bed using existing wireless technologies, Internet of Things (IoT) frameworks, and devices is deployed on an actual production line. The industrial applications targeted in the paper are not time-sensitive and the focus is on framework integration rather than system and application performance. The authors of [4] implemented a rudimentary edge cloud test-bed to quantify the impact of Edge Computing on Mobile Applications using WiFi and the public 4G network. Their effort reveals significant latency and energy usage improvements compared to distant DCs. Additionally, we have in a previous work [5] studied the performance of cloud native applications on commercially available platforms in a smart city context.

iFogSim [6] offers a platform for abstract modeling of resource management techniques in IoT and edge cloud environments, through simulation. In previous works we have also developed several simulators for studying the dynamics of the edge cloud, culminating in [7]. In the absence of a test-bed and in the pursuit of greater flexibility at a higher level of abstraction, simulators are valuable tools. However, we are at the point where a test-bed can be practically implemented and where a simulator cannot capture the complexities of an edge cloud.

Other works attempt to characterise and profile the performance of different aspects of the edge cloud. For example, the authors of [8] evaluate a generic platform for industrial control, with respect to latency, throughput, and CPU load. Their focus is on the pros and cons of virtualisation in a 'Multi-core' environment rather than the 'cloud'. Similarly, [9] implements a water tank control process and evaluates latency over a virtual Software Programmable Logic Controllers (vSoftPLC) on top of LinuxRT. In both of these works, the system implementation rather than the plant under control is evaluated. In [10], the authors propose *Industrial automation as a cloud service*. In the paper the authors evaluate a system of time-sensitive control processes in a 1-tier distributed cloud environment. Latency compensation is modelled and redundancy with stability and smooth controller handover is achieved. True for all systems surveyed above, is that they are neither mission-critical nor time-sensitive at the scale addressed in this paper.

III. RESEARCH TEST-BED

Our research test-bed consists of a 5G radio transmitter and multiple receivers (UEs) (constituting a 5G cell), PC-type compute nodes in DCs and adjacent to the radio transmitter (the edge node), and reduced capacity input-output devices and physical plants at the radio receiver ends. Here, the physical plant can be a mechanical device that continuously performs a task, for example a robotic arm or an autonomous vehicle. We assume that the process which controls the plant is mission-critical and time-sensitive.

We desire a platform with enough knowledge about the application to perform load balancing while allowing an application its own mobility within the network of compute nodes. Further, there is a strong interplay between the edge cloud and the end user equipment such that an application can automatically scale on top of the cloud and provide fall back on local devices. The radio subsystem shall be capable of parallel, synchronized low latency communication with several receivers. We proceed by defining the properties of the system's individual components. From here on the research test-bed is referred to as *the system*.

A. 5G

A 5G wireless system represents the next generation wireless infrastructure [1]. The emerging focus of 5G is URLLC and massive Machine Type Communication (mMTC) where a large number of IoT devices, can reliably be served simultaneously at a low latency, $\leq 5ms$. These conditions cannot be replicated with current 802.11 or Long Term Evolution (LTE) systems.

A next generation wireless network also implies a deeper integration with associated cloud computing resources. On-demand resources are integrated into the Radio Base Stations (RBS) and access networks to off-load the back-haul and eliminate the latency overhead of traversing multiple networks and providers.

Massive Multiple Input Multiple Output (MIMO) is the emerging Radio Access Technology (RAT) for 5G. Fundamentally, massive MIMO is a Multi-User MIMO (MU-MIMO) scheme, which can simultaneously communicate with multiple UEs on the same wireless resource. Additionally, on the RBS-side, massive MIMO operates with significantly more antennas than existing LTE-based RATs. Massive MIMO is typically configured with an order of magnitude more RBS-side antennas than simultaneously served UEs. Consequently, the system's spectral efficiency is a few orders of magnitude greater than existing RATs. The increased spectral efficiency can be used towards serving more simultaneous UEs, increase throughput, or realising mMTC, beyond what can be achieved with existing RATs.

We implement a 5G wireless network using Lund Massive MIMO (LuMaMi). LuMaMi is a Massive MIMO test-bed at Lund University, Sweden. LuMaMi's scope and detailed implementation are found in [11]. LuMaMi can be seen as one solitary 5G cell that can simultaneously communicate with twelve UEs.

LuMaMi is configured according to [12]. This configuration premieres low latency and reliability over high throughput. The modulation scheme is QPSK. The resulting throughput is 4.6 Mbps downlink and 9.1 Mbps uplink, per UE, which is more than sufficient to support our application. LuMaMi allows us to directly route traffic through the system, allowing us to place a compute node in the RBS.

B. Edge cloud and network

The system as a whole is tied together by a set of compute nodes joined by a network. A summary of the compute nodes

Node	Device	Location
Plant	Raspberry Pi 3Bs	Plant adjacent
Edge	Intel Core i7 Desktop	LuMaMi adjacent
ERDC	Intel Core i7 VM	Lund, Sweden.
AWS	Intel Xeon VM	Frankfurt, Germany

TABLE I: Node types

is presented in Table I. The system's network is conceptually configured as depicted in Figure 1. Adjacent to each plant is a Raspberry Pi. In order to sample and manipulate the plant, each Raspberry Pi has been equipped with a ADC/DAC shield. They are compute nodes and may service other functions in addition to interacting with the plant. Each Raspberry Pi is also connected to a 5G UE. The 5G cell is isolated in its own subnet. The subnet includes the wireless infrastructure, an edge cloud node, and plant nodes. The plant-adjacent Raspberry Pis are connected to the system's subnet over LuMaMi. The wireless edge cloud node is adjacent to the LuMaMi RBS. It therefore connects directly to the RBS without traversing additional networks.

A router is connected to the cell's subnet and the larger DCs. The Ericsson Research Data Center (ERDC) resides in Lund, Sweden a few kilometres from the cell. ERDC is a research DC operated by Ericsson (Lund, Sweden), that is open to industrial and academic research efforts within the Wallenberg Autonom Systems and Software Program (WASP). We run on top of Open Stack Pike and our instance (a c4m16) has four Intel i7 cores registered by Linux as 1.6 Ghz, and 16 GB of RAM. The Amazon Web Services (AWS) EC2 instance (a c4.large) is hosted on eu-central-1 (Frankfurt, Germany). This node has two Intel Xeon cores at 2.9 Ghz and 8 GB of RAM. We do not use all cores and therefore expect the latter system to be the best performing. The two Virtual Machines (VMs) on ERDC and AWS connect to the subnet over VPN, allowing direct access between all compute nodes.

C. Cloud native application framework

We define a cloud-native application to be an application that has been disaggregated into logical and independent components connected in a *data-flow graph* and that is hosted on a PaaS framework. The PaaS and its resident applications ubiquitously operate over multiple geographically distributed and heterogeneous compute nodes. An application's data flow graph can be rerouted and extended in run-time, when for example adding a new feature. Additionally, the components shall be able to traverse the cloud and associate with and discover physical input-output devices if the application so requires.

Amazon's AWS, Microsoft's Azure, IBM's Bluemix, and Google's Cloud offer their own flavours of cloud native application platforms, ranging from Software-as-a-Service (SaaS) to server-less Function-as-a-Service (FaaS). However, none of these providers allow their users to define logical data flows nor do they provide necessary performance guarantees. They are typically intended for lifting data from the edge and IoT-devices to the cloud. The services do not provide

native support for closing logical loops from edge devices over the cloud and back to the edge device with guarantees on latency and consistency. For the purpose of building an open research test-bed, these platforms are proprietary and cannot be arbitrarily deployed and independently managed across an edge cloud infrastructure. This is a requirement for our system in order to realize the view where one specification of the software can be deployed anywhere.

In this work, we use Calvin [13] as our cloud platform. Calvin is distributed, event-driven, server-less, and is based on a data-flow programming model. There are a number of such platforms for different workloads, such as: Nebula [14], Node-RED[15], IEC 61499 [16], and Naiad [17]. The aforementioned systems are targeted for the IoT domain and cater for workloads varying from simple event-driven automation to high-throughput Hadoop jobs, but none of them have been built with the intention to run tight control loops over a dynamic distributed system.

Of the above, Calvin is most similar to Node-RED. However, while Node-RED emphasizes the programming model and graphical tools, Calvin puts more focus on runtime dynamics and distributed deployment. The perspective of Calvin is well attuned to the presentation of the Distributed Data-flow model in [18], where a Distributed-NodeRED (DNR) extension is proposed. A notable operational difference is that DNR employs duplication to realize mobility while Calvin's code migration technique is arguably more efficient and is better suited for computationally intense applications [18]. Additionally, in Calvin, we can migrate using various optimisation criteria to provide for instance load balancing or jitter reduction.

Calvin is conceptually structured as follows. The operational units of Calvin are called *actors* (nodes in data-flow) while a *runtime* is an instantiation of the Calvin application environment on a device. In our present implementation there is a one-to-one mapping between Calvin runtimes and compute nodes and we therefore interchangeably refer to them simply as nodes. An actors' input and output messages, are known as *tokens*. A set of actors and their interconnections constitute an *application*. Each node independently schedules its resident actors in a round-robin manner.

Actors' states can be migrated and horizontally scaled across nodes. What constitutes an actors' state is defined by the developer. The Calvin framework can autonomously migrate and place actors to load-balance nodes and to meet its own performance goals. However, application owners can specify requirements for actors which tie them to a preferred runtime. For example, a sensor reading actor can be required to be placed on the node associated with the physical plant it is observing.

IV. EVALUATION

In this section, we present an automatic control application as our time-sensitive mission-critical application. The controller is used to evaluate the performance and plausibility of the test-bed. That is, can we do mission-critical control over

the edge cloud and does the system exhibit the properties associated with an edge cloud. We begin by detailing the application in terms of the software implemented on top of Calvin and the plant that it controls. We then evaluate its performance in a set of experiments designed to:

- 1) Reveal the characteristics of the system and the controller by establishing a baseline observation of the performance and behaviour of the controller over long time periods.
- 2) Verify the adaptability of the system by continuously migrating the controller actor across the system's nodes, in run-time.
- 3) Explore operating limits of the system's nodes and thus their relative advantage by deploying a well-tuned but computationally demanding and time-sensitive controller on the system.

In order to observe the system's performance potential, we choose to limit the study to normal operating conditions. Notably, the connections to the DCs may at times degrade. We assume these to be infrequent, transient behaviours and do not consider how to handle them in this work.

A. Control application

We employ a ball and beam process [19] as our plant under control. The control has to be fast and there are clear limits set by physical constraints, yet enough flexibility for us to modify conditions to create various operating scenarios. Control is critical in that a failure may cause an unrecoverable state.

The objective of the ball and beam process is to expediently move to and maintain a ball on a set location (the set-point) on a beam. The length of the beam is 110 cm. The controller acts on the beam which is manipulated by a motor. The plant *outputs* the angle (α) of the beam and the position of the ball on the beam (x). The control signal, the *input* to the plant, is the radial velocity of the beam (ω). Naturally, the further we go towards the end of the beam the higher the risk that the ball falls off the beam due to network delays, noisy sensors readings, and other system deficiencies.

We implement a controller for the ball and beam process using Model Predictive Control (MPC) [20]. The application periodically samples the position of the ball and the angle of the beam. With every sample the MPC interacts with the plant by changing the velocity of the beam. To figure out what velocity to set the MPC performs a numerical optimisation where it takes into account a series of actions that will bring the ball into the desired state. In a primitive and general form this optimisation may be expressed as

$$\underset{u_0, u_1, \dots}{\text{minimize}} \quad \sum_{t=0}^{T-1} L(x_t, u_t) + \phi(x_T) \quad (1)$$

$$\text{s.t.} \quad x_{t+1} = f(x_t, u_t) \quad (2)$$

$$u_t \in U, x_t \in X \quad (3)$$

where $L(x_t, u_t)$ in Equation (1) is the *cost* function which puts a value to a state x_t and control input u_t at time step t . The function $\phi(x_T)$ assigns a different value specification to

Fig. 2: Calvin MPC implementation

the final (or terminal) state x_T . The number of time steps T is called the horizon and specifies how far into the future the controller anticipates control actions. $f(x_t, u_t)$ in Equation (2) is the plant model which specifies the dynamics of how the system states evolve with the time step t . Equation (3) is a set of expressions which state limitations to the plant inputs and the state space. All of this is defined to set the operating conditions for the controller. The initial state, x_0 , is drawn from measurements and state estimation. The end result is a quadratic program which is re-evaluated every sample.

For the optimisation, we use a dynamically linked binary created with the use of QPgen [21]. In our implementation, we use a sampling period of 50 ms (20 Hz). This choice is a reasonable trade-off between control performance and early observations of the system's latencies. Unlike the well known and often used Proportional Integral Derivative (PID) controller, the execution of an MPC is demanding and the execution time is not constant. The time it takes to solve its optimisation routine varies with disturbances acting on the system and where within its operating range it is currently acting. QPgen is an efficient solver and our optimisation problem has few variables, yet we shall see that the time it takes to find a solution can be considerable. Sometimes there is no solution or one is very hard to find. In such a case we end the search after a fixed amount of iterations in the optimiser.

The processor and memory demands of the MPC optimization may be considerable and the many ways of tuning it for various situations make it interesting as an edge cloud application. A number of controllers can be designed for the same problem where computational and memory demands are weighted to performance, operational range, stability regions, and erratic behaviour. To handle the presence of plant and sensory noise we include a Kalman filter. The filter is also used to estimate the speed of the ball. We use a simple plant,

a basic MPC controller and a standard state estimator but even our rather simple case allows us to study and demonstrate behaviour in the experimental platform and the effects on the control.

The Calvin application graph for the MPC control loop is shown in Figure 2. The rounded rectangles represent individual components, implemented as actors, which are deployed onto the systems. The two Analog to Digital Converter (ADC) sensory actors adhere to component reuse and the principle idea that they need not be collocated. However, they are to be read jointly and therefore share a clock tick. The components within the grey area can be freely placed within the system. The ADCs (the position and angle sensors) and the Digital to Analog Converter (DAC) (the motor actuator) have an affinity to the plant-adjacent node.

We note that the scheduling and inter-node communication in the software platform introduce a significant amount of delay and jitter, which affects the performance of the controller and cause oscillations. In extension, much can be done in terms of model and controller tuning, state estimation, delay prediction, system improvements etc., but in this work we are interested in studying the overall performance of the platform. Remedies and improvements are left for later work, here we characterise the basic system and do not focus on details of control performance.

All related software runs on top of Linux and the Calvin runtime is launched using real-time priority and the POSIX FIFO scheduling policy. We allow the edge node to take full advantage of this as it runs on bare metal. The kernels have not been patched with the PREEMPT_RT patch set [22].

B. System characteristics

To characterise and verify the basic functionality of the system we ran our MPC on each of the nodes in Section III-B. With each test we let the MPC control the beam for 60 minutes while alternating the set-point of the ball between the centre position and one side of the beam. To be robust in this experiment, we restrict the set-point with a large margin to the end of the beam. On the other hand, the further out we move the ball the more we put the controller to work, which is something we return to in Section IV-D.

Figure 3a shows the Round-Trip Times (RTTs) from the Raspberry Pi at the plant to the other systems. Notably the wireless link realised with LuMaMi introduces a latency of 5ms one way, as made evident by the 10ms RTT between the plant and the edge node. 5G is pushing for even faster RTT but this is good radio link performance compared to commercially available alternatives.

Figure 3b shows the MPC execution time. We chose a simple scenario where all nodes can execute the MPC with a significant margin. However, clearly, the Raspberry Pi at the plant is many times slower than the other systems. The AWS node is faster than edge and ERDC which is to be expected by the specification in III-B. We see in this graph that there are large outliers in terms of execution time at the plant. Even though we execute in real-time mode we have seen recurrent

Fig. 3: Statistical summary of the MPC baseline measurements.

extended system interruptions to the MPC on the Raspberry Pi. This could be the cause of these outliers. On the DCs we do not expect real-time properties to apply outside the virtual machine and can therefore expect some outliers. Due to the short execution times, we expect them to be unlikely.

Figure 3c shows the aggregate latency from reading the position of the ball to applying the control signal (i.e. adjusting the velocity of the beam). This is an important measure because the controller is designed with the assumption that the input to output is instantaneous and hence, that the state of the system has not changed when the control signal is applied. We see that the differences in delay are not as pronounced as in Figure 3a. The execution times in Figure 3b and the network latency in Figure 3a are not the only contributors to the control latency. This tells us that a significant proportion of the delay in our system is introduced by the software platform or application design and not the network. The system dynamics causing the increasing variance in Figure 3c is also an interesting topic for further research. The effect on our process due to these properties are visualised in Figure 3d where we see that the energy of the control signal u increases the further we move from the plant. Notice that the AWS node performs well but network delays causes it to exhibit a larger mean and variance in the control signal. Such an effect can be part of the heuristics when deciding where to place control in the edge cloud.

C. System adaptability

We have established that a controller can successfully be implemented on the edge cloud test-bed and studied characteristics in terms of execution times, latencies and jitter. Essential to the mutability of the system is its ability to migrate applications and actors to respond to the applications' and the infrastructure's changing objectives. During a migration the Calvin cluster performs the necessary modification of the network communication path, recreates the actor at the target node, copies state and handles the transition of token queues. Although the actor moves point-to-point, changing the communication paths may involve many nodes in the cluster.

We now add this perspective by way of relocating the MPC amongst the nodes while balancing and repositioning the ball as in Section IV-B.

In Figure 4 we continuously migrate the MPC actor randomly across the four compute nodes, in run-time. When doing this, the system must ensure to keep the Kalman filter, the set-point, the previous state, and tracing meta data intact. Delays, data loss or duplication, and incorrect state transfer negatively impacts the control performance. Figure 4c shows the placement of the actor in time. Figure 4b and Figure 4a show the controller inputs and outputs respectively.

Figure 4b shows that the process is stable and is able to operate without interruptions. The ball stays on the beam and close to the desired position. Figure 4a confirms what is presented in Figure 3d, i.e., the control signal increases as a function of the distance to the plant. Set-point changes are clearly visible has high peaks but the migrations in Figure 4c are not evident in Figure 4a nor Figure 4b. However, the peak in the control signal near the set-point change after 650 seconds is likely caused by a coinciding migration. In its current form, the system is not aware of when or to where a migration will occur nor do we attempted to mitigate its potential effects.

D. Tightened constraints

We now demonstrate an example use case where the controller takes advantage of the edge cloud. Here, relative to the previous example, we have set a constraint on the control signal to the plant. Albeit being a synthetic exercise, it is not an unreasonable action since limits in control signal are commonly used to, for instance, reduce actuator wear and to avoid non-linear parts of the operating range. To make the associated optimisation problem harder we move the ball just short of the end of the beam. In combination with the constraints this will cause a higher load on the MPC host node. Small disturbances may cause the ball to fall off.

The experiment shown in Figure 5 applies this configuration. The graphs show time series of the inputs to the controller, the execution time of the MPC, and the total latency

Fig. 4: Time-series of MPC being randomly migration between the system's four nodes.

from input to output of the measured nodes. The blue circles mark occasions when the MPC fails to find a solution.

As constraints are tightened it becomes increasingly hard to find a control signal sequence which moves the system from the present state to the target state, while staying within the bounds. This manifests in longer execution times for the optimisation. As seen in the execution times in Figure 5, when the system has settled around a set-point, the optimisation is easy to solve and computationally light. In these situations the controller on the plant performs well. Latency and jitter of the networked controllers may cause them to deviate more from the set-point.

Eventually however, the computational limitations of the plant cause the ball to falls off the beam as a set-point change occurs. Note that the plant is not unable to move the ball to the position at the end of the beam but eventually, model errors and noise become too large for it to handle. In contrast, the edge node is able to operate without failure. Also note that on the AWS instance, despite its computational capacity, the controller fails to cope with the resulting latency and system jitter - the ball falls off.

The execution time at the plant when the MPC fails to find a feasible solution, is close to an order of magnitude that of the sampling time, represented by a line which extends well beyond the top of the graph. This is representative of the computational problems experienced at the plant due to noise during a set-point change. In contrast, an equal number of iterations consumes 80 ms on the edge node and only 40 ms on the AWS.

With the position closer to the end of the beam and with reduced range in the control output signal, the controller repeatedly experiences non-trivial situations which require additional iterations of the optimisation loop. At times, noisy readings make a tough situation even worse and there may seemingly be no solution that keeps the ball on the beam.

When a new evaluation can be made quickly enough then the state of the system may still be such that the ball can be saved. On the AWS node the communication delays increase the frequency of these tough situations and we see repeated problems of finding a solution. However, the speed of the AWS node allows it to cope with many of these situations since the combined execution and communication delay is much less than the compute time at the plant. Only our edge node is in a position where it is able to handle the full range of the system noise.

V. CONCLUSIONS

In this paper we presented an edge cloud research test-bed for an IoT and heterogeneous cloud environment. We deployed an automatic control application on the test-bed to act on a time-sensitive and mission-critical process. We evaluated the controller's viability, performance, and system characteristics. Our evaluation shows that cloud native control loops can viably be deployed on the edge cloud. We operate at a sampling rate of 20 Hz but there is potential for this to be pushed further in the near future. We also showed that our controller can benefit from the edge cloud and that the system and the placement of the controller can be dynamically reconfigured in run-time without strictly sacrificing stability.

To further improve the software platform we see that there remain work to be done on actor scheduling, overheads of message passing, and synchronisation of such a tightly connected system in an inherently uncertain cloud environment. With improvements, we can increase the requirements on the application and the system to move towards anticipated future applications which require such a system to be user friendly, self-adaptive, resilient, high performing, and deliver low latency and low jitter.

Importantly, we conclude that we have an observable test-bed which enables us to continuously operate mission-critical

Fig. 5: Time-series of experiments run with tightened constraints on the plant, egde, and AWS nodes.

applications while performing targeted experiments. We have now arrived at a system which allows broad experimental research of the interplay between the application and the underlying platform. Continuations to the work targets novel and established techniques within the fields of control theory, distributed systems, and software engineering.

The project source code is available on GitHub [23].

ACKNOWLEDGEMENTS

We would like to thank Andreas Johansson and Martin Nilsson at the Department of Electrical and Information Technology at Lund University for their help with designing and constructing the Raspberry Pi shields needed for our experiments. This work is funded by the Swedish Research Council (VR) under contract number C0590801, the Lund Center for Control of Complex Engineering Systems (LCCC) also funded by VR, and by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

REFERENCES

[1] M. Shafi, A. F. Molisch, P. J. Smith, T. Haustein, P. Zhu, P. De Silva, F. Tufvesson, A. Benjebbour, and G. Wunder, “5g: A tutorial overview of standards, trials, challenges, deployment, and practice,” *IEEE Journal on Selected Areas in Communications*, vol. 35, no. 6, pp. 1201–1221, 2017.

[2] J.-M. Kang, H. Bannazadeh, and A. Leon-Garcia, “Savi testbed: Control and management of converged virtual ict resources,” in *2013 IFIP/IEEE International Symposium on Integrated Network Management (IM 2013)*. IEEE, 2013, pp. 664–667.

[3] J. Wan, S. Tang, Z. Shu, D. Li, S. Wang, M. Imran, and A. V. Vasilakos, “Software-defined industrial internet of things in the context of industry 4.0,” *IEEE Sensors Journal*, vol. 16, no. 20, pp. 7373–7380, Oct 2016.

[4] W. Hu, Y. Gao, K. Ha, J. Wang, B. Amos, Z. Chen, P. Pillai, and M. Satyanarayanan, “Quantifying the impact of edge computing on mobile applications,” in *Proceedings of the 7th ACM SIGOPS Asia-Pacific Workshop on Systems*, ser. APSys ’16. New York, NY, USA: ACM, 2016, pp. 5:1–5:8. [Online]. Available: <http://doi.acm.org/10.1145/2967360.2967369>

[5] W. Tärneberg, V. Chandrasekaran, and M. Humphrey, “Experiences creating a framework for smart traffic control using aws iot,” in *Proceedings of the 9th International Conference on Utility and Cloud Computing*, ser. UCC ’16. New York, NY, USA: ACM, 2016, pp. 63–69. [Online]. Available: <http://doi.acm.org/10.1145/2996890.2996911>

[6] H. Gupta, A. Vahid Dastjerdi, S. K. Ghosh, and R. Buyya, “ifogsim: A toolkit for modeling and simulation of resource management techniques in the internet of things, edge and fog computing environments,” *Software: Practice and Experience*, vol. 47, no. 9, pp. 1275–1296, 2017.

[7] W. Tärneberg, A. Papadopoulos, A. Mehta, J. Tordsson, and M. Kihl, “Distributed approach to the holistic resource management of a mobile cloud network,” in *1st International Conference on Fog and Edge Computing IC FEC 17, 14 May 2017, Madrid, Spain*, 2017.

[8] N. Mahmud, K. Sandström, and A. Vulgarakis, “Evaluating industrial applicability of virtualization on a distributed multicore platform,” in *Emerging Technology and Factory Automation (ETFA), 2014 IEEE*. IEEE, 2014, pp. 1–8.

[9] C. Horn and J. Krüger, “Feasibility of connecting machinery and robots to industrial control services in the cloud,” in *2016 IEEE 21st International Conference on Emerging Technologies and Factory Automation (ETFA)*. IEEE, 2016, pp. 1–4.

[10] T. Hegazy and M. Hefeeda, “Industrial automation as a cloud service,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 26, no. 10, pp. 2750–2763, 2015.

[11] S. Malkowsky, J. Vieira, L. Liu, P. Harris, K. Nieman, N. Kundargi, I. C. Wong, F. Tufvesson, V. Öwall, and O. Edfors, “The world’s first real-time testbed for Massive MIMO: Design, implementation, and validation,” *IEEE Access*, vol. 5, pp. 9073–9088, 2017.

[12] W. Tärneberg, M. Karaca, A. Robertsson, F. Tufvesson, and M. Kihl, “Utilizing massive mimo for the tactile internet: Advantages and trade-offs,” in *SECON Workshops - Robotic Wireless Networks*. IEEE, 2017.

[13] P. Persson and O. Angelsmark, “Calvin—merging cloud and iot,” *Procedia Computer Science*, vol. 52, pp. 210–217, 2015.

[14] M. Ryden, K. Oh, A. Chandra, and J. Weissman, “Nebula: Distributed edge cloud for data intensive computing,” in *2014 IEEE International Conference on Cloud Engineering (IC2E)*. IEEE, 2014, pp. 57–66.

[15] Node-red. [Online]. Available: <https://nodered.org>

[16] V. Vyatkin, “Iec 61499 as enabler of distributed and intelligent automation: State-of-the-art review,” *IEEE Transactions on Industrial Informatics*, vol. 7, no. 4, pp. 768–781, 2011.

[17] D. G. Murray, F. McSherry, R. Isaacs, M. Isard, P. Barham, and M. Abadi, “Naiad: a timely dataflow system,” in *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*. ACM, 2013, pp. 439–455.

[18] N. K. Giang, M. Blackstock, R. Lea, and V. C. M. Leung, “Developing iot applications in the fog: A distributed dataflow approach,” in *2015 5th International Conference on the Internet of Things (IOT)*, Oct 2015, pp. 155–162.

[19] M. Virseda, “Modeling and control of the ball and beam process,” *MSc Theses*, 2004.

[20] J. Rawlings and D. Mayne, *Model Predictive Control: Theory and Design*. Nob Hill Pub., 2009. [Online]. Available: https://books.google.se/books?id=3_rfQQAACAAJ

[21] P. Giselsson, “Qpgen,” <http://www.control.lth.se/user/pontus.giselsson/qpgen/>, 2018.

[22] L. Foundation, “Real-time linux,” <https://wiki.linuxfoundation.org/realtime/start>, 2018.

[23] P. Skarin and W. Tärneberg, “Control over the cloud software,” <https://github.com/pskarin/CotC>, 2018.